UG Exam Performance Feedback
Third Year
2018/2019 Semester 2

COMP34412 Natural Language Systems

Goran Nenadic
Tingting Mu

Comments

Question 1. The average mark for this question was 55%. Part (i) had no problems and the majority of students got full marks. Part (ii) expected the students to talk about specific challenges for machine translation, and most students had provided good discussions. For part (iii), there have been many cases where the idea of TBL wasn't properly explained, in particular confusing the training and application parts (e.g. what it means to correct POS errors); many marks are lost here. Part (iv) was generally reasonably well answered, but there were marks lost for not describing what sequence labeling is, or for not discussing all possible features. Finally, part (v) aimed to demonstrate different approaches, but many answers failed to generalize the issues from this particular example to a more generic problem.

Question 2. The average mark for this question was 68 – so the majority of students did very well. Part (i) was bookwork, and while the majority of answers replicated examples from the class, there were some with really excellent examples designed by students; still, there have been few cases where nominal coordination wasn't explained. Part (ii) was in most cases answered in full, and the majority of students did the example in part (iii) correctly; there were some confusion in what right/left arc means – it's nothing to do with drawing the lines to the left/right; still this wasn't a major issue – only few answers were wrong. Part (iv) was reasonably answered, but with many answers been generic and with no taking the given case study in the focus; many failed to discuss how to segment the feedback text and how to align it to a specific course unit; some answers discussed for example relationship extraction without specifying what kind of relations are needed.

Question 3 was answered in generally well. Most students understood the basic principal of Naïve Bayes classifier, and Part i.b was answered well. However, some students were not able to apply the knowledge to analyse the examples given in the question, and lost marks. In Part i.c, most students understood the principal of simplified Lesk, but some students only explained the principal without applying it to analyse the example given in the question. Almost all the students did well in Part ii, demonstrating good understanding on vowel and consonant pronunciation.

Question 4 was answered in generally well. Most students did well in Part i.a, b and d, demonstrating good understanding of term-document matrix and term similarity. Many students did not do well in Part i.c), and they were not able to apply Skip-Gram knowledge to analyse the given example in the question. Part ii was not answered well, and the understanding on feature extraction from speech signals could have been improved.